The Challenge of Data Scale

Dan Reed
reed@renci.org

Chancellor's Eminent Professor
Vice Chancellor for Information Technology
University of North Carolina at Chapel Hill

Director, Renaissance Computing Institute (RENCI)

renci
renaissance computing institute

You Might Be A Big System Geek If …

- **You think a $2M cluster**
  - is a nice, single user development platform
- **You need binoculars**
  - to see the other end of your machine room
- **You measure system network connectivity**
  - in hundreds of kilometers of cable/fiber
- **You dream about cooling systems**
  - and wonder when fluorinert will make a comeback
- **You telephone the local nuclear power plant**
  - before you boot your system
- ***You order storage systems***
  - *and remember when a gigabyte was lots of data*

renci
renaissance computing institute

The Fujitsu Eagle: Remember?

- **Features**
  - *380 MB capacity* and 18 ms access time
  - data rate of nearly 2 MB/s
  - 130 pounds and 10” platters
  - ~$10K list

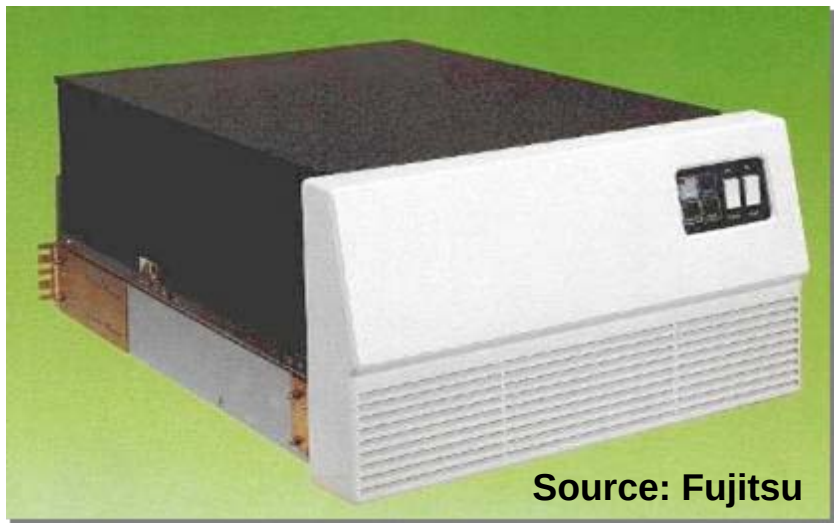

renci
renaissance computing institute

Presentation Outline

- **Complex application examples**
  - scaling and interdisciplinary coupling
  - sensor data explosion
- **Technology enablers**
  - exponential change and implications
  - lessons from other sectors
    - biomedicine and business
- **Meditations on the future**
  - alternative data models and tools
  - the proverbial box
    - what does it look like on the outside?

renci
renaissance computing institute

21st Century Scaling Challenges

- **Population growth**
  - severe weather sensitivity
    - national impact
  - geobiology and environment
  - economics and finance
  - sociology and policy

Consilience

- **Economics and healthcare**
  - longitudinal public health data
    - environmental interactions
  - genetic susceptibility
    - heart disease, cancer, Alzheimer's
  - privacy and insurance
  - public policy and coordination

DNAdirect Your Genes In Context

renci
renaissance computing institute

Weather and Economic Loss

- **$10T U.S. economy**
  - 40% is adversely affected by weather and climate
- **$1M in loss to evacuate each mile of coastline**
  - we now over warn by 3X!
  - average over warning
    - 200 miles, or $200M per event
- **Improved forecasts**
  - lives saved and reduced cost
- **LEAD national Grid**
  - Oklahoma, Indiana, UCAR
  - Colorado State, Howard, Alabama
  - Millersville, NCSA, North Carolina

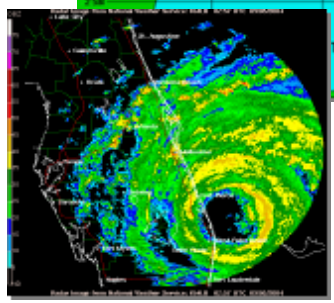

Source: Kelvin Droegemeier, Oklahoma

renci
renaissance computing institute

The LEAD Vision: A Paradigm Shift

Analysis and Assimilation

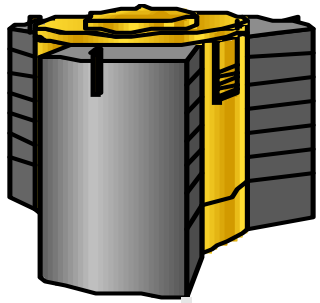

Prediction and Detection

Product Generation and Dissemination

Models/Algorithms Driving Sensors

The challenge: Build cyberinfrastructure services that provide adaptability, scalability, availability, usability, and real-time response.

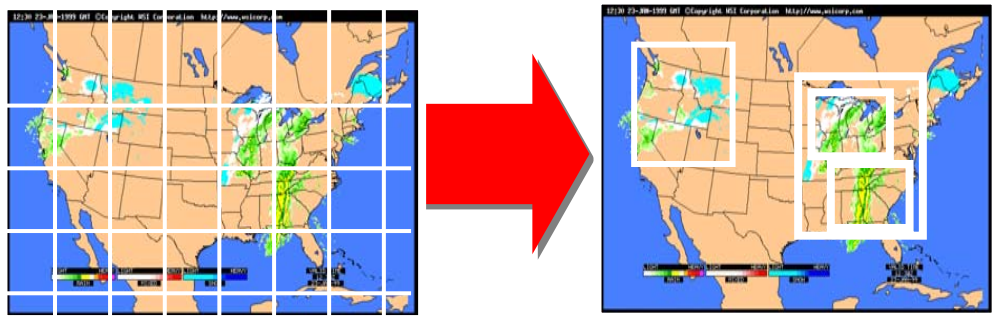

End Users

Source: Plale, Indiana

renci
renaissance computing institute

Genetics and Disease Susceptibility

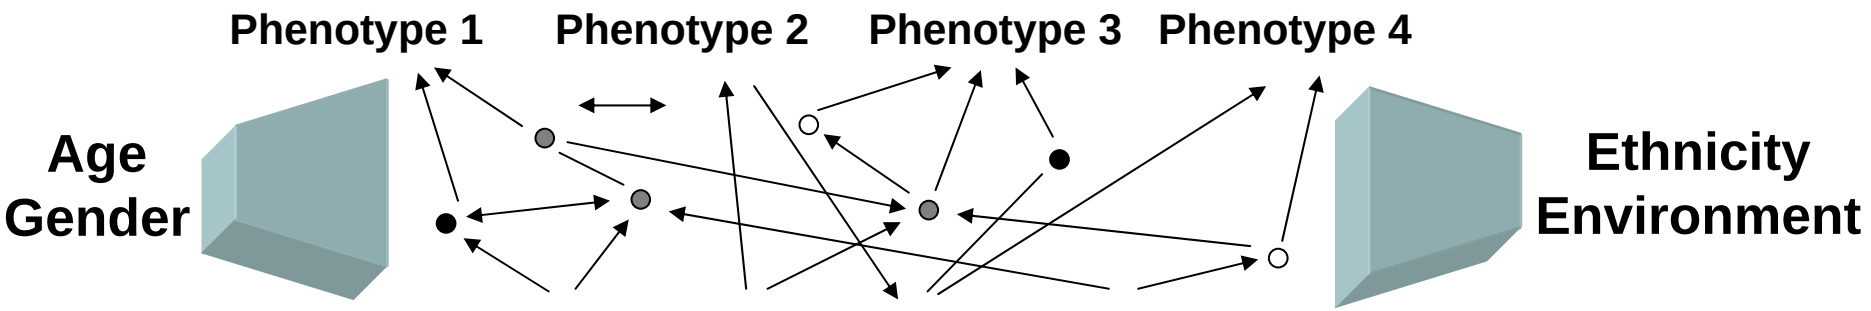

Identify Genes

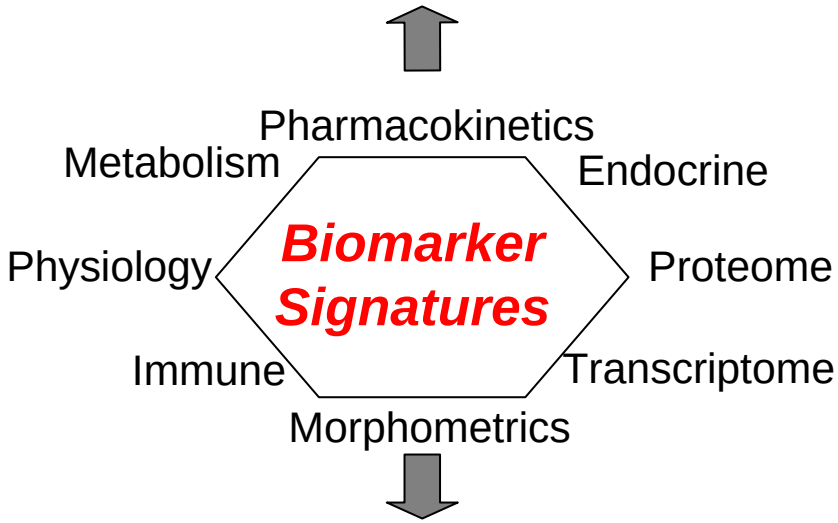

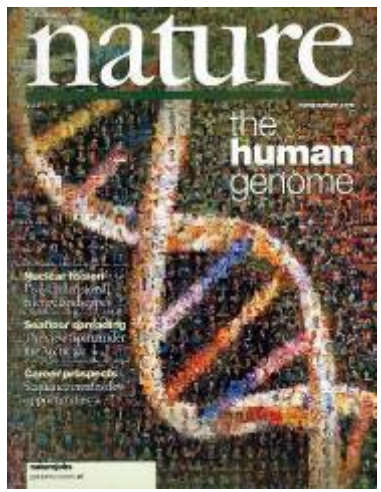

Predictive Disease Susceptibility

Source: Terry Magnuson, UNC

renci
renaissance computing institute

Mean Onset of Alzheimer's Disease

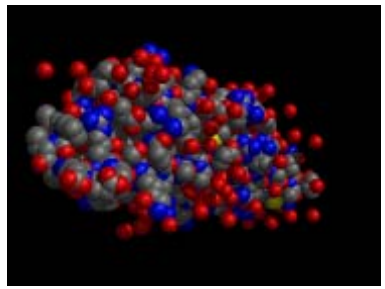

- **apolipoprotein (apo)**
  - apoE2, apoE3 and apoE4 alleles
    - on chromosome 19
  - apoE4 allele
    - 40% to 60% of Alzheimer's patients
- ***Not* the only Alzheimer's cause**

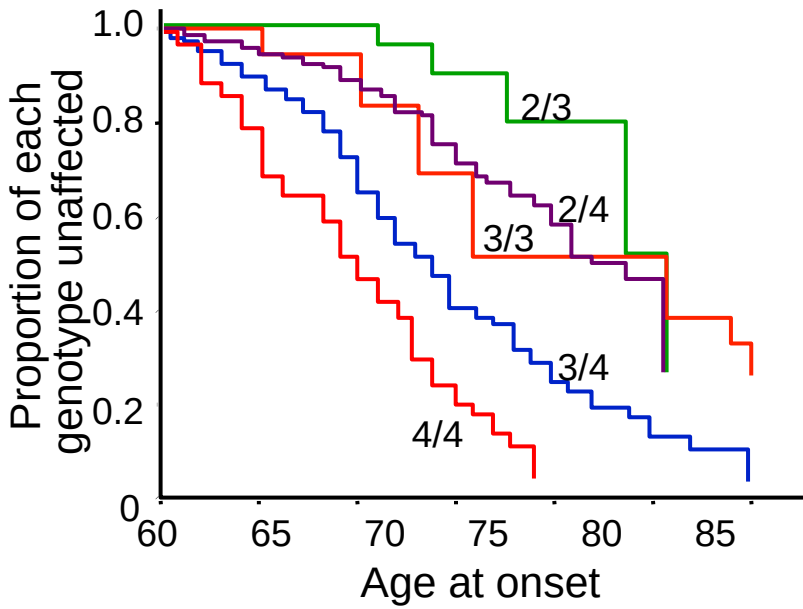

- **apo gene inheritance**
  - ~25% inherit 1 copy of apoE4 allele
    - Alzheimer's risk increases 4X
  - *2% inherit 2 copies of apoE4 allele*
    - *Alzheimer's risk increases 10X*

Source: Alan Roses, GSK

renci
renaissance computing institute

The Data Tsunami

- **Many sources**
  - agricultural
  - biomedical
  - environmental
  - engineering
  - manufacturing
  - financial
  - social and policy
  - historical
- **Many causes and enablers**
  - increased detector resolution
  - increased storage capability
- **The challenge: *extracting insight!***

We Are Here!

renci
renaissance computing institute

PITAC: Data Management and Sensors

- **Explosive growth in the resolution of sensors and scientific instruments is creating unprecedented volumes of experimental data.**
- **We must increase investment and focus on sensor- and data-intensive computational science in recognition of the explosive growth of experimental data, itself a consequence of increased computing capability.**

renci
renaissance computing institute

Sensor Data Overload

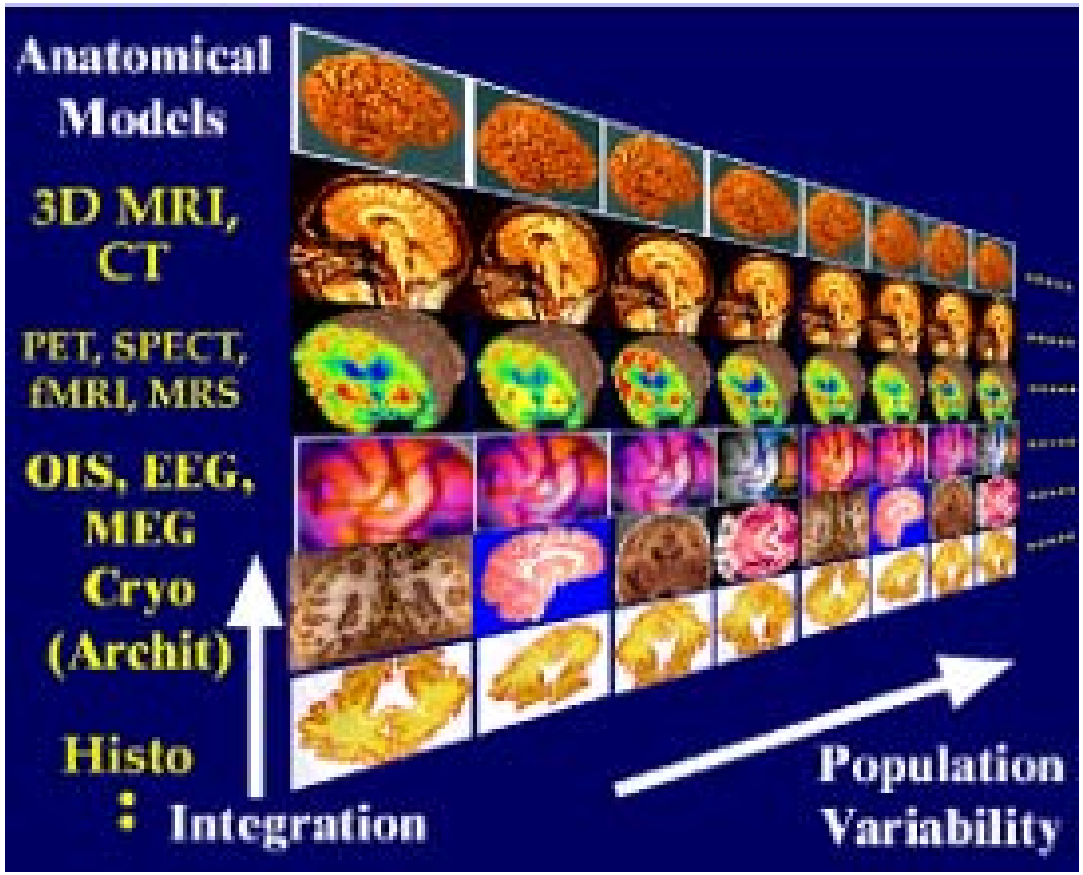

Source: Chris Johnson, Utah/Art Toga, UCLA

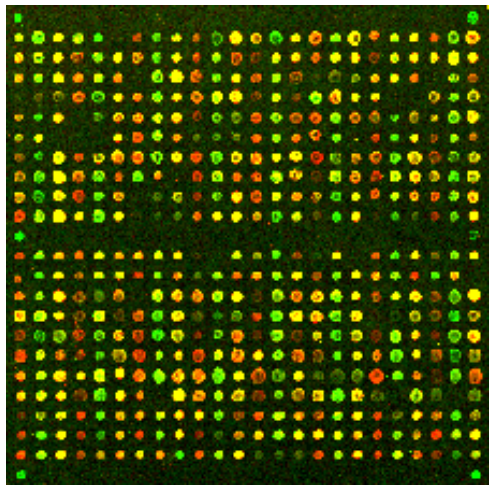

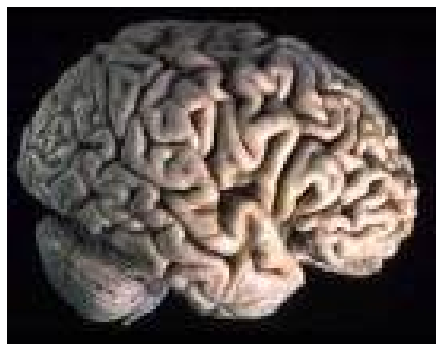

AFFYMETRIX GENECHIP®

- **High resolution brain imaging**
  - 4.5 petabytes (PB) per brain

renci
renaissance computing institute

Digital Reality: The Exponentials

1956

1972

2006

- **Megabyte**
  - a small novel
- **Gigabyte**
  - a pickup truck filled with paper or a DVD
- **Terabyte: one thousand gigabytes – ~$1000 today**
  - the text in one million books
  - entire U.S. Library of Congress is ~ten terabytes of text
- **Petabyte: one thousand terabytes**
  - 1-2 petabytes equals all academic research library holdings
    - coming soon to a pocket near you!
  - *soon routinely generated annually by many scientific instruments*
- **Exabyte: one thousand petabytes**
  - 5 exabytes of words spoken in the history of humanity

Source: Hal Varian, UC-Berkeley

renci
renaissance computing institute

Data Growth and Complexity

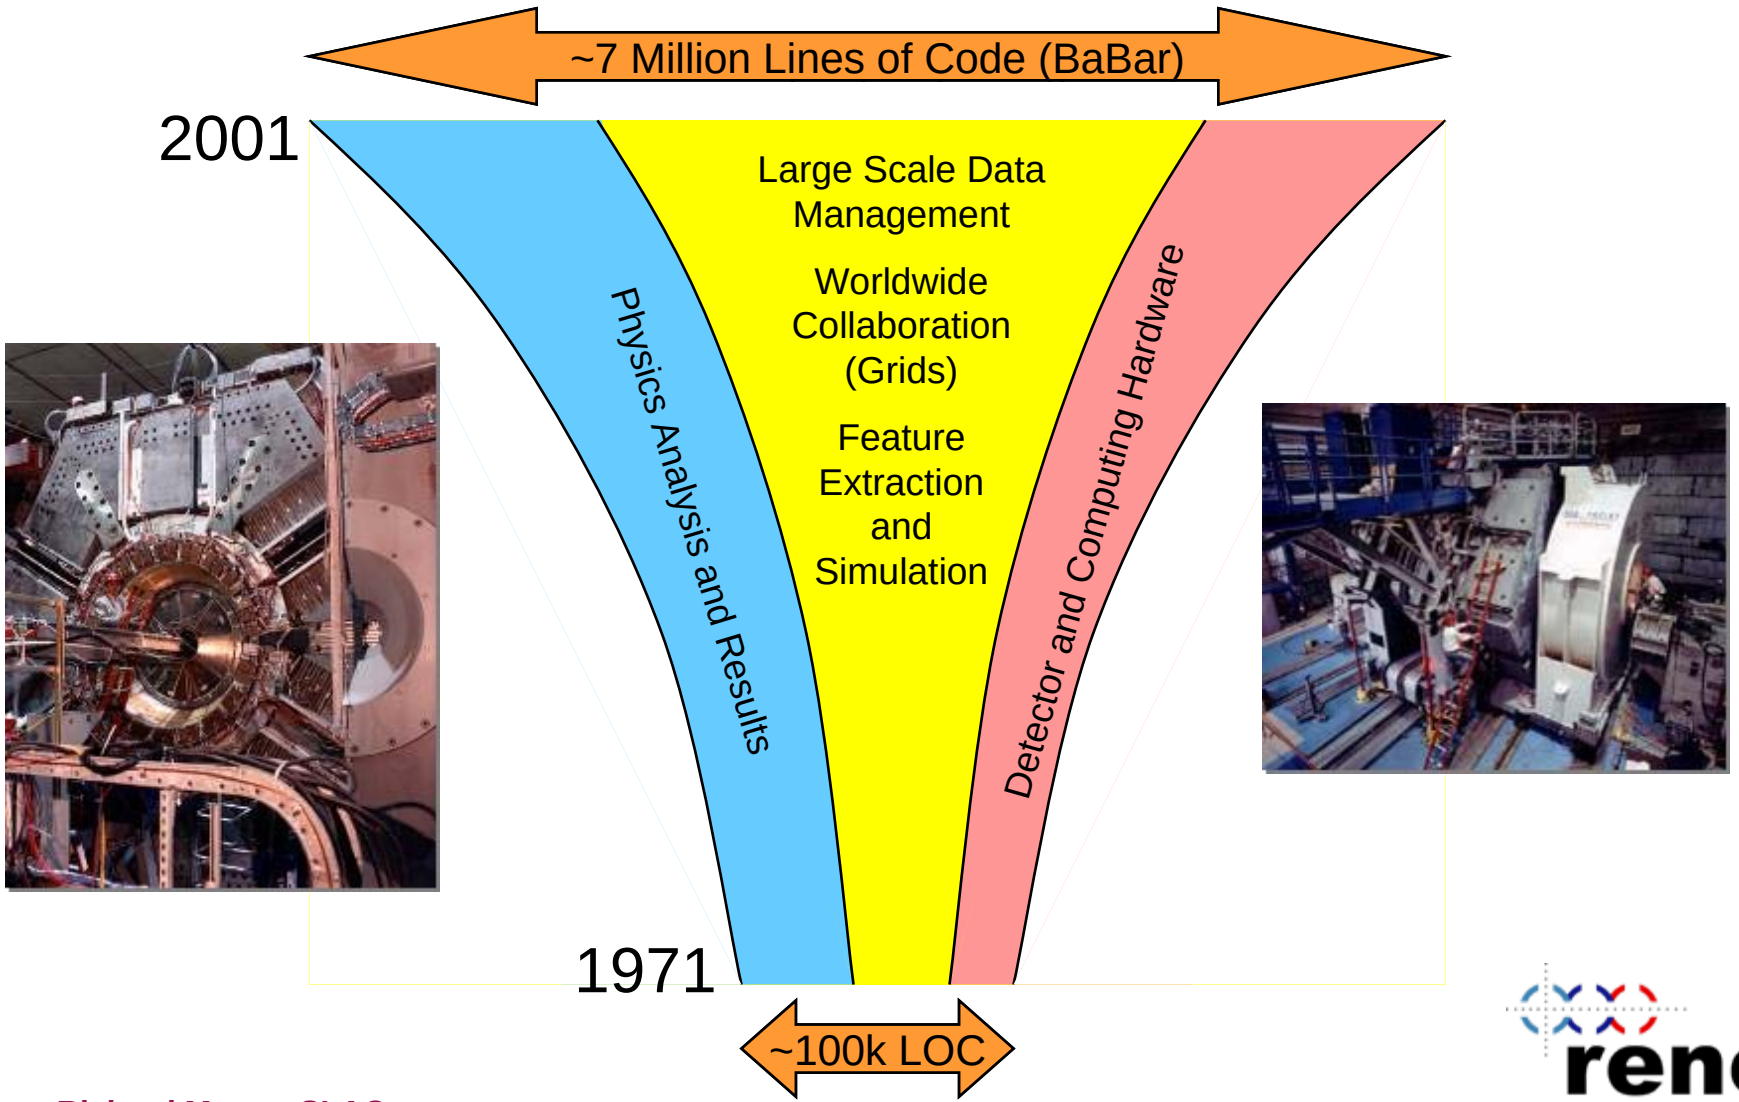

Source: Richard Mount, SLAC

renci
renaissance computing institute

Consumer Gigapixel Images

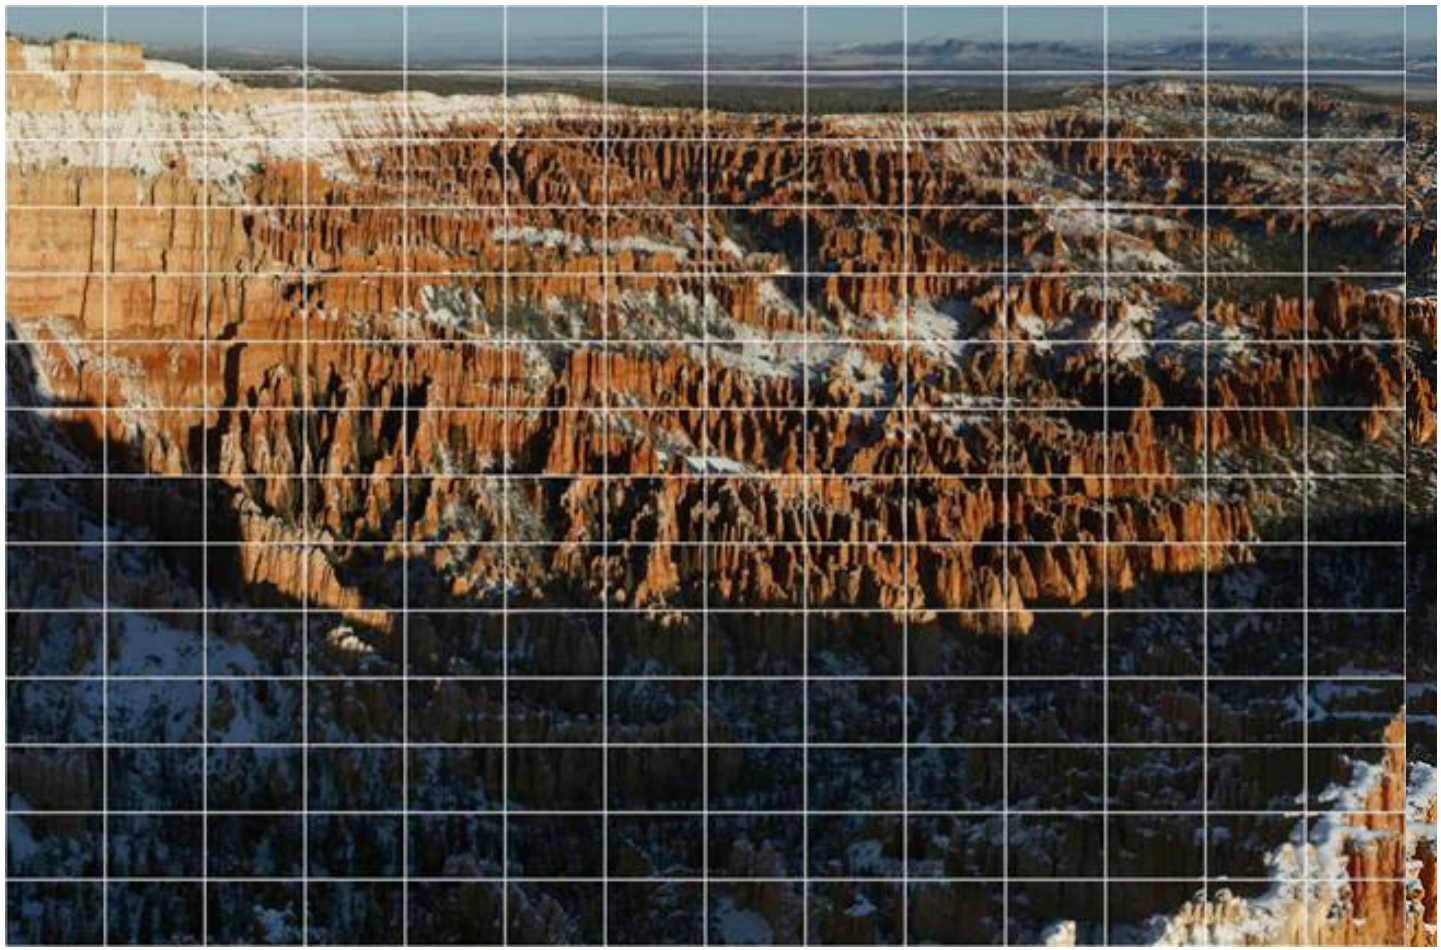

- **40,784 x 26,800 pixels**
  - 196, 6 megapixel images

Source: Max Lyons, www.tawbaware.com/maxlyons/gigapixel.htm

renci
renaissance computing institute

Astronomy and Consumer Cameras

- **Digital camera sales**
  - now exceed analog
  - January 2006
    - Nikon stopped film camera production
- **From glass plates to CCDs**
  - detectors follow Moore's law
    - consumer electronics
  - data tsunami
    - data doubles every two years
- **Telescope growth**
  - 30X glass (concentration)
  - 3000X in pixels (resolution)
- **Single astronomy images**
  - 16Kx16K pixels and growing

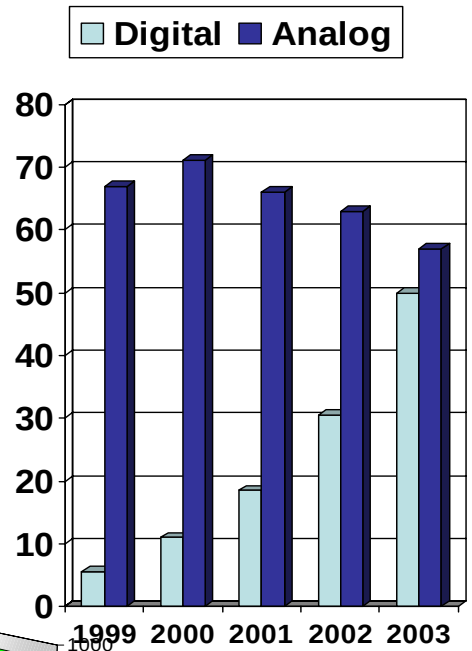

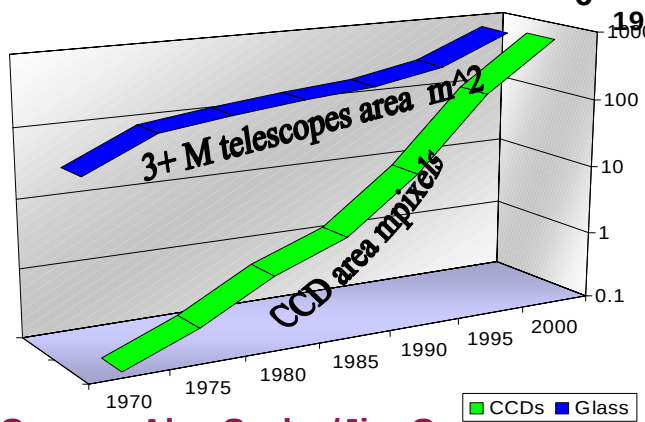

Source: Alex Szalay/Jim Gray

renci
renaissance computing institute

Large Synoptic Survey Telescope (LSST)

- **Top project of the astronomy decadal survey**
- **Celestial cinematography**
  - 3 gigapixel detector for wide field imaging
- **Science**
  - beyond the standard model
    - non-baryonic dark matter
    - non-zero Λ and neutrino oscillations
  - observation targets
    - near Earth object survey
    - weak lensing of wide fields
    - supernovae measurements
- **Features**
  - 9.6 square degree field/6.5 meter effective aperture
  - *~15 TB of data/night, target first light 2012*

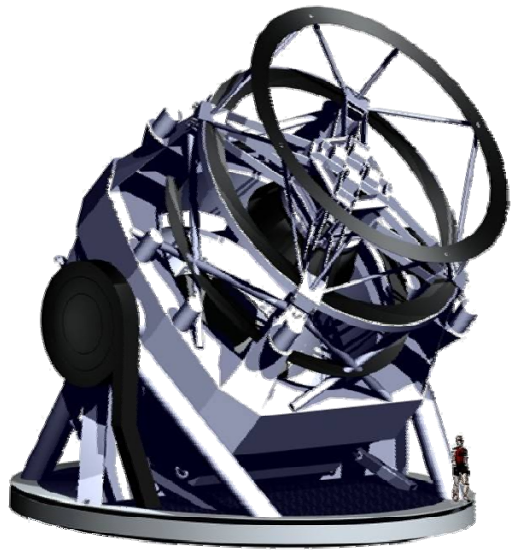

renci
renaissance computing institute

Lessons from Astronomy

- **Historically, discoveries accrued to those**
  - with access to unique data
  - who built next generation telescopes
- **Two things changed**
  - growing costs and complexity of telescopes
  - emergence of whole sky surveys
- **The result – virtual astronomy**
  - discovering significant patterns
    - analysis of rich image/catalog databases
  - understanding complex astrophysical systems
    - integrated data/large numerical simulations

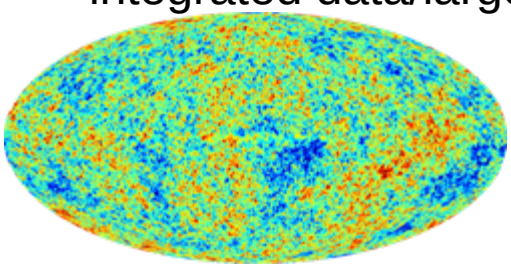

SDSS

2MASS
2 MICRON ALL SKY SURVEY

NVO
NATIONAL VIRTUAL OBSERVATORY

renci
renaissance computing institute

{Inter}national Virtual Observatory

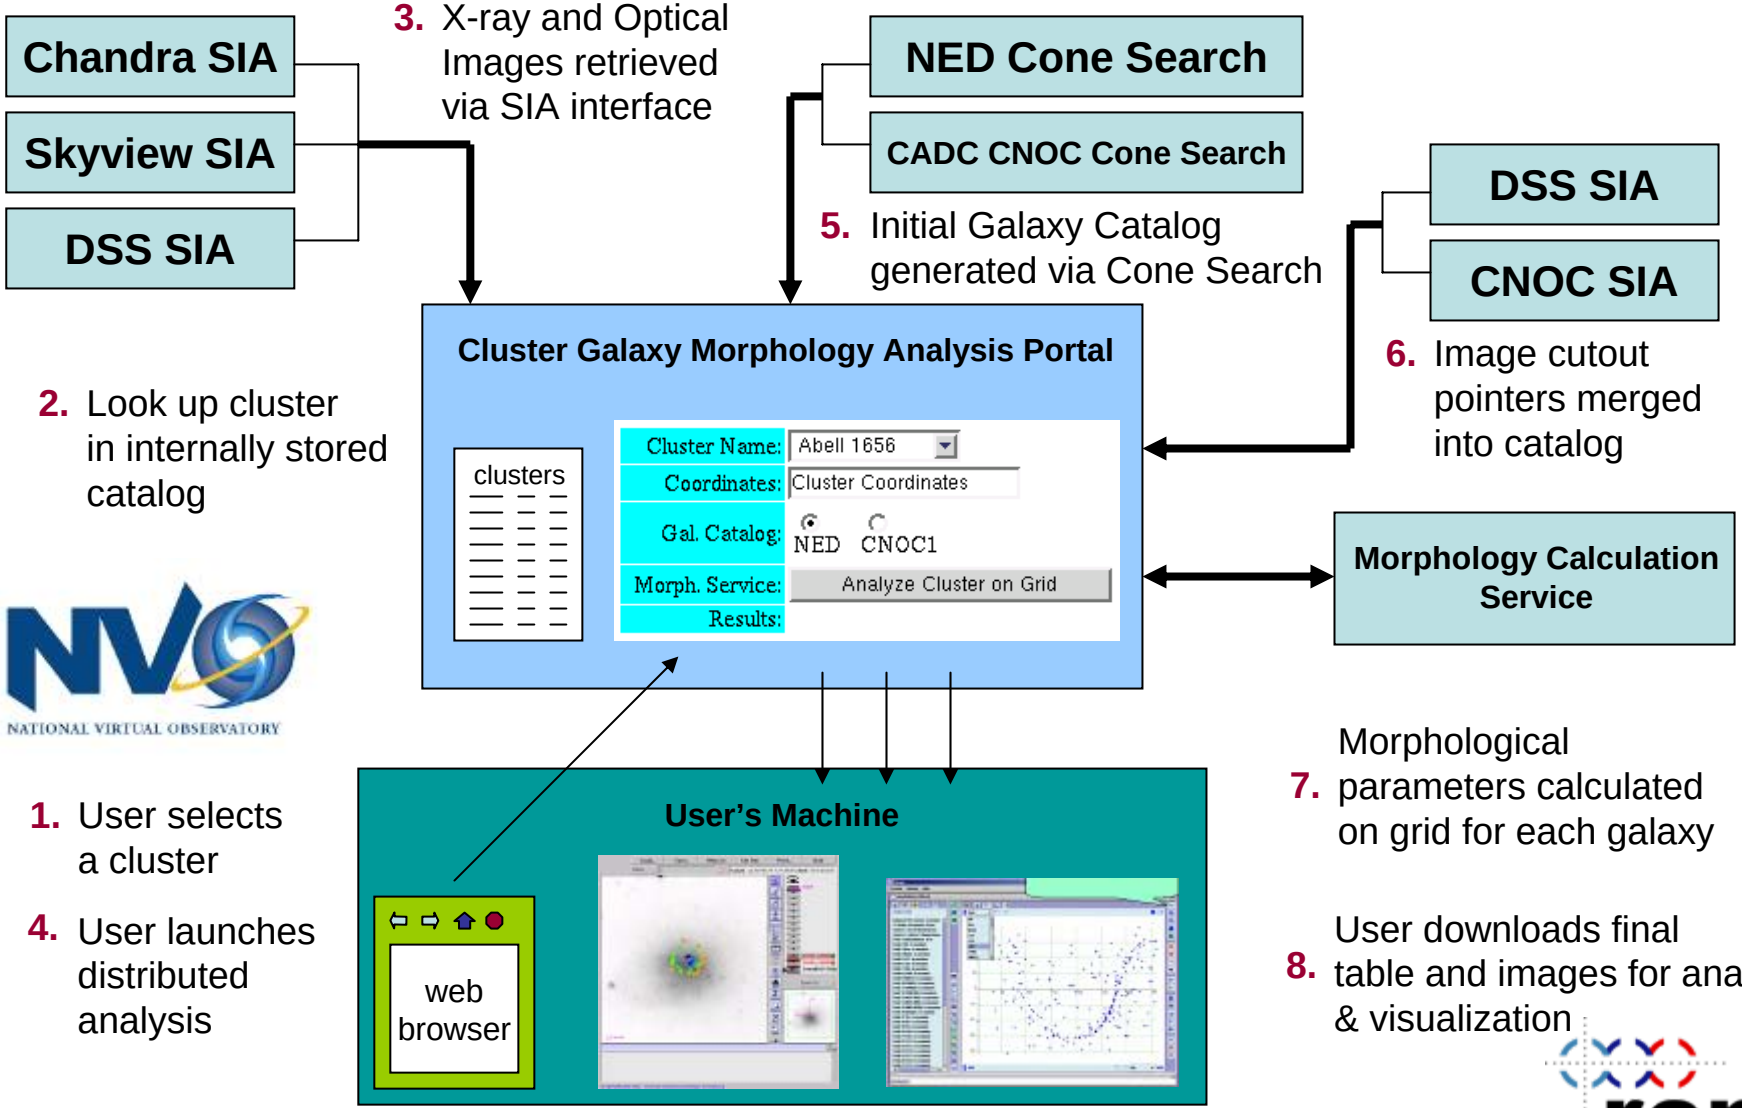

Source: Ray Plante, NCSA

renci
renaissance computing institute

Gene Expression and Microarrays

- **Concurrent evaluation**
  - expression levels for thousands of genes
- **Photolithography**
  - up to 500K 10-20 micron cells
    - each containing millions of identical DNA molecules
- **Image capture and analysis**
  - laser scanning and intensity calculation

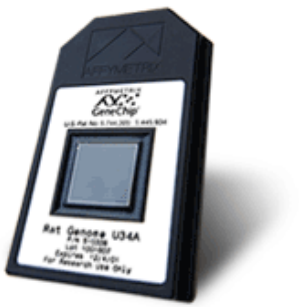

AFFYMETRIX GENECHIP®

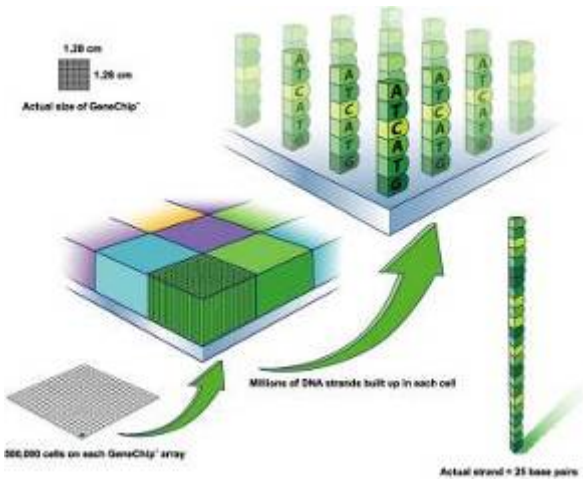

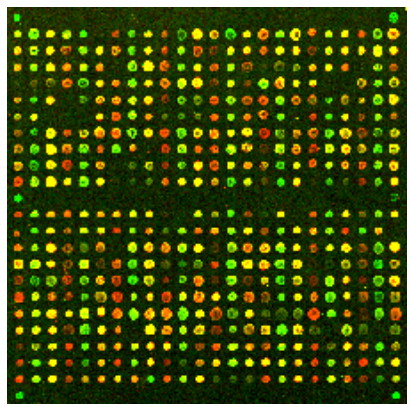

Source: Affymetrix

renci
renaissance computing institute

Data Heterogeneity and Complexity

Genomic, proteomic, transcriptomic, metabalomic, protein-protein interactions, regulatory bio-networks, alignments, disease, patterns and motifs, protein structure, protein classifications, specialist proteins (enzymes, receptors)

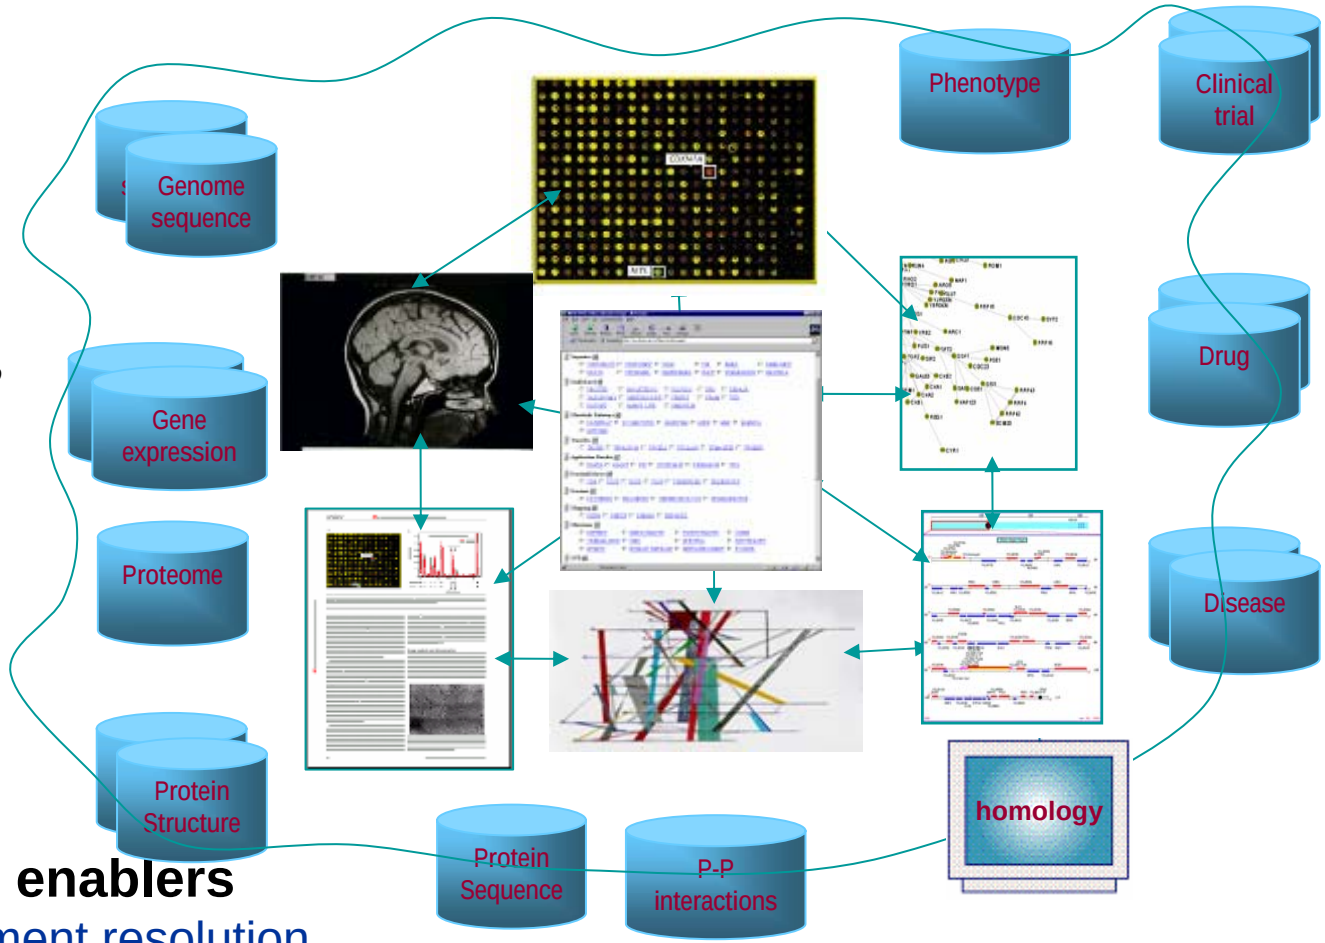

- **Many causes and enablers**
  - increased instrument resolution
  - increased storage capability
- **The challenge: *extracting insight!***

Source: Carole Goble, Manchester

renci
renaissance computing institute

Need: Simple, Easy-To-Use Tools

“Genome. Bought the book. Hard to read.”

Eric Lander

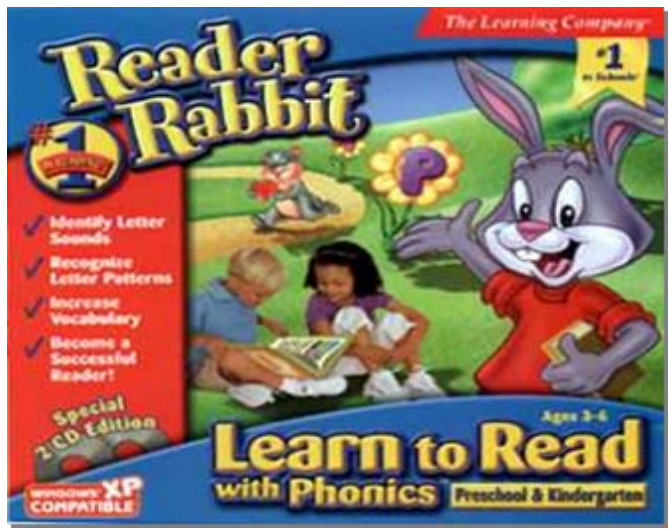

renci
renaissance computing institute

The Problem for the e-Scientist

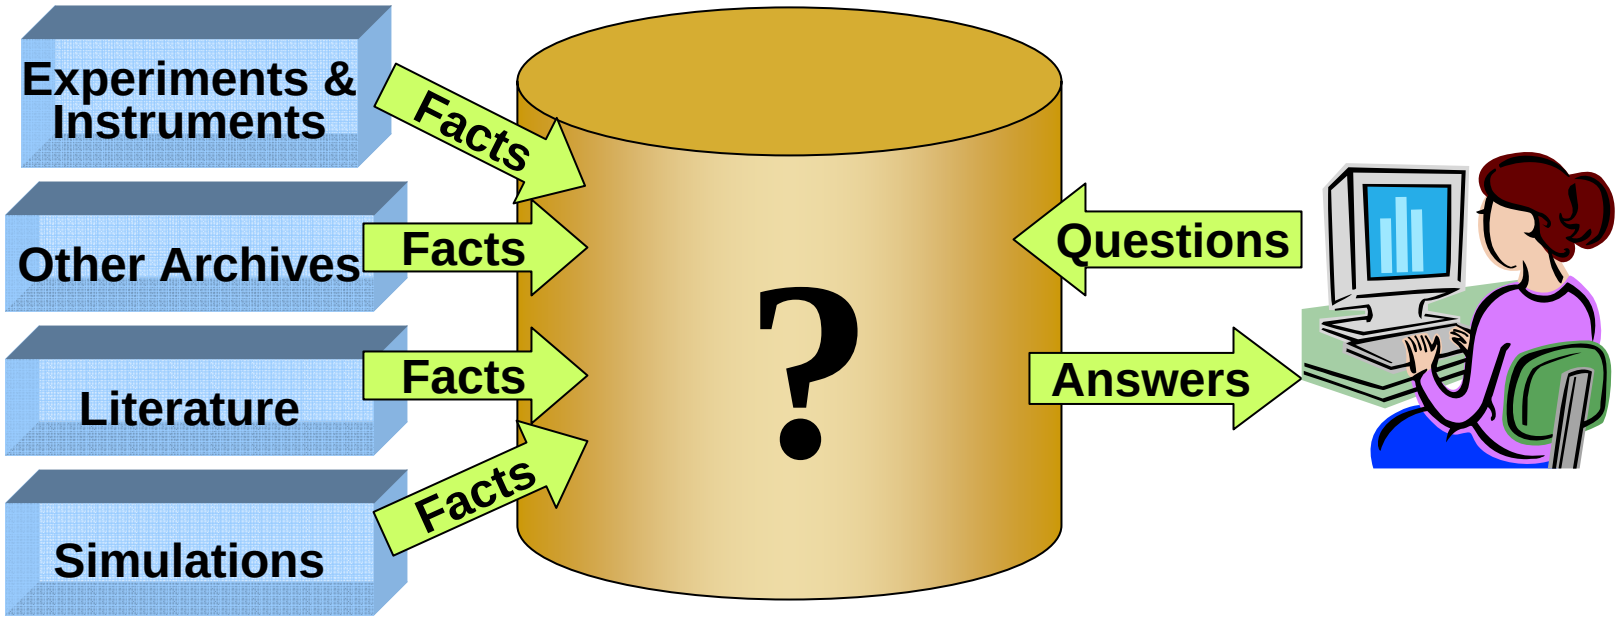

- Data ingest
- Managing a petabyte
- Common schema
- How to organize it?
- How to *re*organize it?

- Query and visualization tools
- Support and training
- Performance
  - execute queries in a minute
  - batch (big) query scheduling

Source: Tony Hey, Microsoft

renci
renaissance computing institute

It's Called Google …

"I want a file system visible from earth orbit."

Mark Seager, LLNL

… but the file system isn't as important now

renci
renaissance computing institute

Two Data Analysis Modes

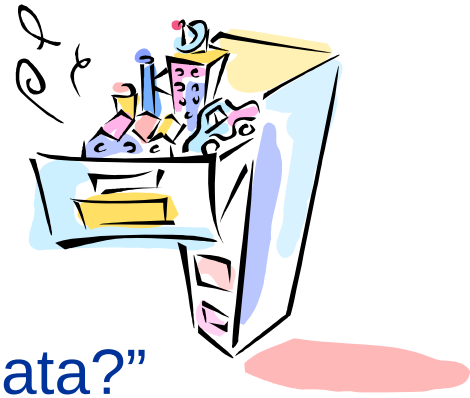

- **Hypothesis-driven**
  - "I have an idea, let me verify it."
- **Exploratory**
  - "What correlations can I glean from the data?"
- **Different tools and techniques**
  - exploratory analysis relies on deep data mining
    - supervised and unsupervised learning
  - "grep" is not a data mining tool
- **Distributed, multidisciplinary data**
  - rising rapidly and correlations needed

renci
renaissance computing institute

Data Mining: Why We Care

- **An iterative process**
  - discovering valid, novel and useful patterns
  - valid: generalize to the future
  - novel: what we do not know
  - useful: actionable based on insight

What information consumes is rather obvious: it consumes the attention of its recipients. Hence *a wealth of information creates a poverty of attention*, and a need to allocate that attention efficiently among the overabundance of information sources that might consume it.

Herbert Simon

renci
renaissance computing institute

The Wal-Mart Lesson

"We know how many 2.4-ounce tubes of toothpaste sold yesterday, and what was sold with them. Our database grows because we capture data on every item, for every customer, for every store, every day" **Dan Phillips, VP of IS, Wal-Mart**

WAL★MART®
ALWAYS LOW PRICES.
Always.
Walmart★com

- **~3600 U.S. stores and ~500 TB**
  - more importantly, they react immediately

renci
renaissance computing institute

From RFID to Smart Dust

- **Smart dust (pixie dust)**
  - wireless environmental sensors
    - perhaps as small as 1 mm^3
  - commodity hardware and MEMS
  - flora and fauna measurements
    - an IP address for every frog
- **RFID tags**
  - secure, inexpensive and disposable
  - passive and active versions
  - contents: identity, state, location
  - logistics management and tracking
    - Wal-Mart leadership and EU Euro tracking

UCB COTS Smart Dust

SUN SPOT

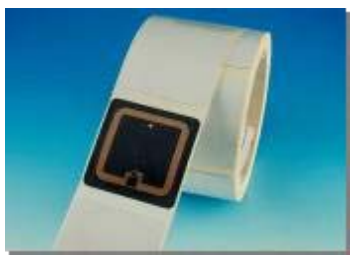

renci
renaissance computing institute

RENCI Petascale Data System

- **Not just a large parallel file system!**
  - rather, a database and data model
    - ontologies, ingestion, federation, mining, …
- **Two research targets**
  - science and humanities domains
    - biomedicine, environmental analysis, economics
  - large-scale data management technology
- **Rationale**
  - exploratory analysis, not just hypothesis-driven studies
  - diverse communities with high-level interfaces
- **grep is not a search methodology!**

renci
renaissance computing institute

From SNPs to HapMap

- **Single Nucleotide Polymorphisms (SNPs)**
  - one in 1200 bases differ across individuals
  - SNPs act as markers to locate genes
- **Common groups of SNPs are shared**
  - i.e., form a haplotype
- **HapMap data sources**
  - 90 Yoruba individuals (30 trios) from Nigeria (YRI)
  - 90 individuals (30 trios) of European descent from Utah (CEU)
  - 45 Han Chinese individuals from Beijing (CHB)
  - 45 Japanese individuals from Tokyo (JPT)
- **3,500,000 SNPs typed**
  - basis for association studies for disease identification

renci
renaissance computing institute

UNC HapMap Simulator/Bakeoff

- **Resample from HapMap haplotypes**
  - create individuals with statistical properties of data
  - recombine and adjust
    - biased SNP selection and sample size
- **Model disease**
  - create large populations with families/select individuals
  - disease model can be complex
    - involving multiple loci
- **Enable analysis bakeoff**
  - five data sets simulated with 500K SNPs
    - trait caused by common sequence variants
    - each data set has 5000 cases/5000 controls
      - common versus rare traits
      - independent versus additive versus epistatic
      - variation in effect size and allele frequency
  - blind analysis by five UNC groups
    - computer science, applied math
    - biostatistics, pharmacy and genetics

renci
renaissance computing institute

Structuring Scientific Literature

- **Each domain described by distinct sources**
  - ontologies
    - domain concepts, entities and relationships
  - taxonomies
    - classification into ordered categories
  - curated databases
    - additional inter-entity links/relations

- **Scenario**
  - customize multidimensional space
    - change source/target ontology
  - visualize
    - connections (concepts/entities, chain/paths)
    - titles/abstracts of documents supporting the connections

Source: Anita Lungu, Duke

renci
renaissance computing institute

Federation and Information Visualization

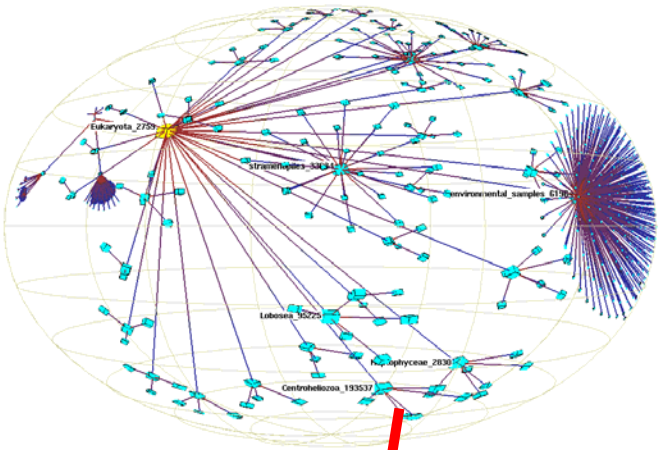

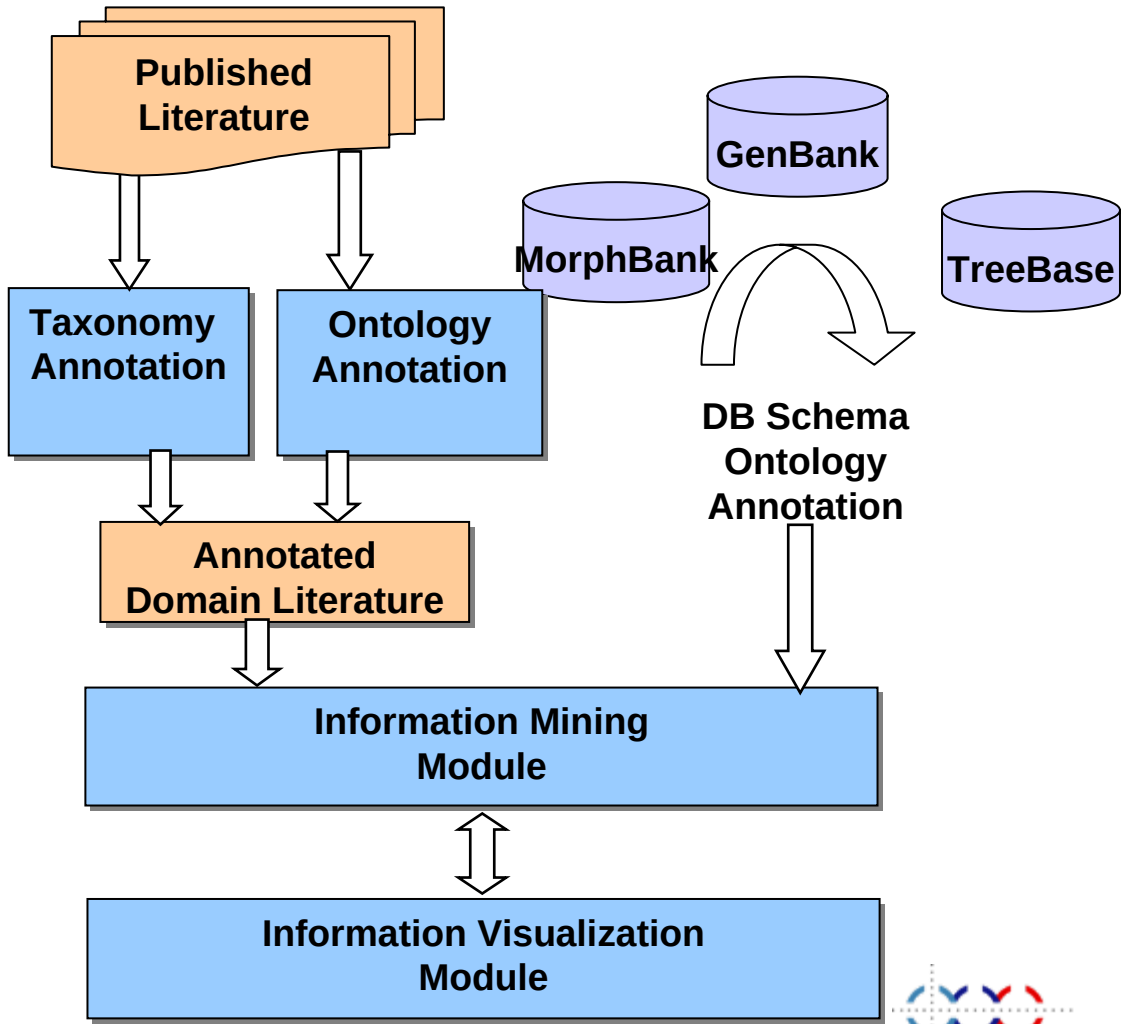

Source: Anita Lungu, Duke

renci
renaissance computing institute

Memex: Still Prescient

"Consider a future device for individual use, which is a sort of mechanized private file and library. It needs a name, and to coin one at random, "memex" will do. A memex is a device in which an individual stores all his books, records, and communications, and which is mechanized so that it may be consulted with exceeding speed and flexibility. It is an enlarged intimate supplement to his memory."

Vannevar Bush

"As We May Think," 1945

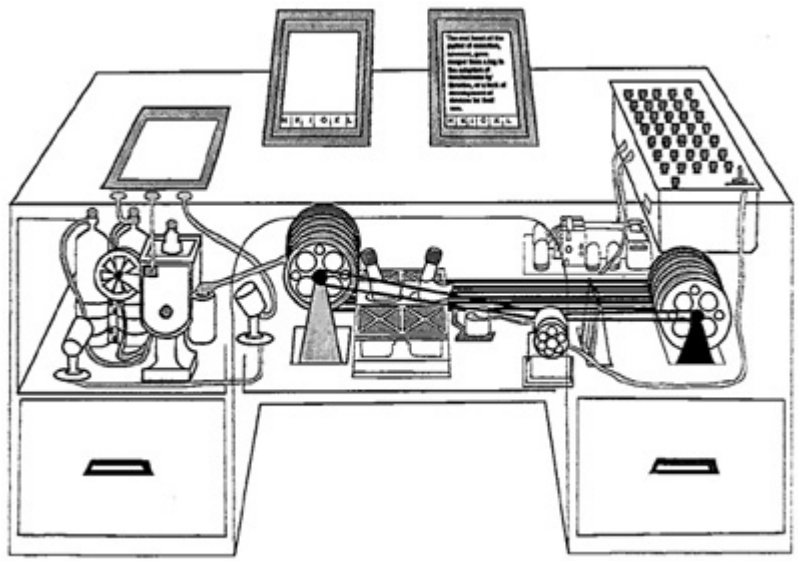

renci
renaissance computing institute